

# Understanding Personal Concept in Open-Vocabulary Semantic Segmentation

Sunghyun Park\* Jungsoo Lee\* Shubhankar Borse Munawar Hayat  
 Sungha Choi† Kyuwoong Hwang Fatih Porikli  
 Qualcomm AI Research‡

{sunpar, jungsool, sborse, hayat, sunghac, kyuwoong, fporikli}@qti.qualcomm.com

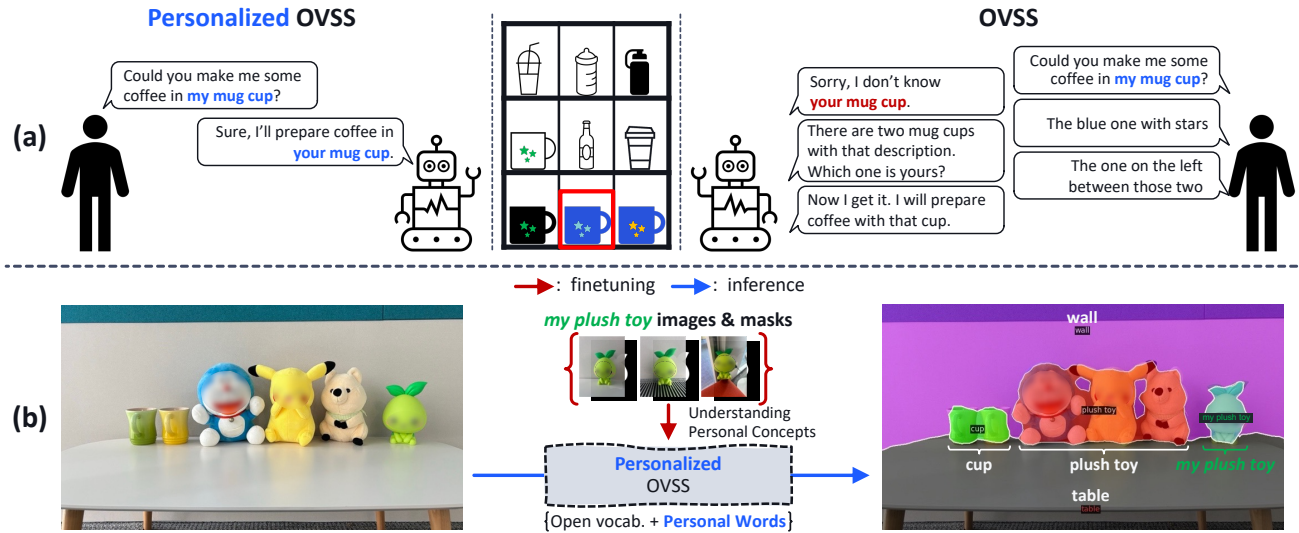


Figure 1. Description of personalized open-vocabulary semantic segmentation (personalized OVSS). (a) While existing OVSS models require detailed descriptions to find ‘my mug cup’, personalized OVSS enables to find it directly. (b) Given a few images and masks of my personal concept (e.g. my plush toy), our proposed framework finetunes an existing OVSS model to distinguish the personal concept among objects within the same class (e.g. my plush toy among multiple plush toys).

## Abstract

While open-vocabulary semantic segmentation (OVSS) can segment an image into semantic regions based on arbitrarily given text descriptions even for classes unseen during training, it fails to understand personal texts (e.g. ‘my mug cup’) for segmenting regions of specific interest to users. This paper addresses challenges like recognizing ‘my mug cup’ among ‘multiple mug cups’. To overcome this challenge, we introduce a novel task termed personalized open-vocabulary semantic segmentation and propose a text prompt tuning-based plug-in method designed to recognize personal visual concepts using a few pairs of images and masks, while maintaining the performance of the original OVSS. Based on the observation that reducing false predictions is essential when applying text prompt tuning

to this task, our proposed method employs ‘negative mask proposal’ that captures visual concepts other than the personalized concept. We further improve the performance by enriching the representation of text prompts by injecting visual embeddings of the personal concept into them. This approach enhances personalized OVSS without compromising the original OVSS performance. We demonstrate the superiority of our method on our newly established benchmarks for this task, including  $FSS^{per}$ ,  $CUB^{per}$ , and  $ADE^{per}$ .

## 1. Introduction

Due to the recent advances in large-scale vision-language models (e.g. CLIP [32]), open-vocabulary semantic segmentation (OVSS) [37, 38], the task of performing segmentation with a set of arbitrary open-vocabulary texts, has recently shown large improvements. Unlike traditional semantic segmentation [5] that is limited to making segmen-

\*These two authors contributed equally to this work. †Project lead.  
 ‡Qualcomm AI Research is an initiative of Qualcomm Technologies, Inc.




Figure 2. How personalized open-vocabulary semantic segmentation (personalized OVSS) differentiates with existing tasks. Few-shot segmentation utilizes a few images and masks for capturing a certain visual concept (e.g. pokeball), without considering making predictions with an arbitrary set of words. While OVSS predicts objects with an arbitrary set of vocabulary (e.g. table and orange), it fails to recognize a given visual concept (e.g. pokeball). Unlike the two tasks, personalized OVSS recognizes the personal concept along with its original OVSS predictions. Additional to the novel task, we also propose a plug-in method applicable to existing off-the-shelf OVSS methods.

tation predictions within a fixed set of categories, OVSS enables to segment regions with arbitrary classes that are not used during the training phase. Such models are crucial for deploying semantic segmentation models in real-world applications (e.g. robot assistants) since we may encounter novel objects that were unseen during training.

Despite the recent remarkable improvements in OVSS [8, 25, 37–40, 43], segmenting the region based on user’s interests using unseen categories has been underexplored. For example, as shown in Fig. 1 (a), recognizing ‘my mug cup’ among a number of mug cups in a cupboard is challenging for existing OVSS studies. The main reason is that they are designed to distinguish between mug cups and other classes (e.g. bottles and plastic cups) rather than identifying a certain mug cup among several mug cups. This feature is important when considering its application in real-world scenarios. For example, with current OVSS models, we have to provide detailed descriptions to ask our robot assistant to identify ‘my mug cup’, which is cumbersome. On the other hand, personalized OVSS enables the robot assistants to identify ‘my mug cup’ even without such detailed descriptions.

Although there is another group of studies that focuses on learning certain visual concepts using few images and masks, referred to as few-shot segmentation [11, 15, 24], these studies fall short in two key aspects for this particular task. First, they are not designed to perform open-vocabulary segmentation, which is crucial for real-world applications that need to recognize a wide range of objects. Second, they do not consider identifying the personal visual concept among objects sharing the same category (e.g. finding ‘my plush toy’ among other plush toys). As described in Fig. 2, we propose a new task termed *personalized open-vocabulary semantic segmentation*, which aims to segment regions of user’s interest by using an arbitrary set of words and learning personal visual concepts with few images and masks. The main goals of this task are to 1) recognize the given personal concept and 2) maintain the original OVSS performance. While Yo’LLaVa [29] re-

cently focuses on personalized visual question answering using large multimodal models, learning personal concepts along with open-vocabulary segmentation models trained on large-scale datasets has been underexplored, despite its importance in practical applications requiring personalized object recognition.

In this paper, along with our newly proposed task, we propose a *simple yet effective* method that can be applied to off-the-shelf pretrained OVSS models. We first utilize text prompt tuning, which focuses on learning the newly given visual concept using an additional learnable textual embedding. While such an approach indeed recognizes the personal visual concept at a reasonable level, it is accompanied by an increased number of false positive predictions. To address such an issue, we propose a ‘negative mask proposal’ that adds a mask to focus on learning the regions other than the personalized concept. Through the negative mask proposal, we prevent the model from making overconfident predictions during text prompt tuning. Additionally, we further improve the performance by injecting the visual embeddings extracted from a pretrained image encoder (e.g. CLIP [32]) into the textual embedding.

The contributions of our work are as follows:

- To the best of our knowledge, this is the first work to propose the task termed *personalized open-vocabulary semantic segmentation*, which segments images with both an arbitrary set of words and the given personal visual concept included in a few pairs of images and masks.
- We propose a *simple yet effective* plug-in method applicable to existing open-vocabulary segmentation models, which performs text prompt tuning with visual information injected and a negative mask proposal.
- We achieve state-of-the-art performance on three newly established personalized open-vocabulary segmentation benchmarks, including  $FSS^{per}$ ,  $CUB^{per}$ , and  $ADE^{per}$ .

## 2. Related Work

**Open-Vocabulary Semantic Segmentation.** In order to enable segmentation models to segment images with unseen

categories, previous open-vocabulary segmentation models [1, 3, 9, 37] have concentrated their efforts on mapping the visual features to the semantic space. Recently, the enhanced capacity of cross-modal alignment (*e.g.* CLIP [32]) has led to a proliferation of studies utilizing pretrained vision-language models for OVSS [8, 10, 14, 20, 31, 38–40, 43]. Side Adapter Network (SAN) [40] attaches a lightweight side network to a pretrained CLIP, which focuses on predicting mask proposals while leveraging the features of the original CLIP. Also, based on the observation that the representation space of the Diffusion model [33] is semantically differentiated and localized well, Open-vocabulary Diffusion-based panoptic SEgmentation (ODISE) [38] leverages the image encoder of diffusion models for segmenting images into open-vocabulary labels. Although such models show superb performance on OVSS, they exhibit shortcomings in segmenting regions of user’s interest. To address these challenges, this paper proposes a plug-in method applicable to off-the-shelf OVSS models, enabling the learning of new visual concepts.

**Personalization of Semantic Segmentation.** Performing semantic segmentation for the objects of specific interests included in one or a few exemplary images, referred to as personalization of semantic segmentation or few-shot segmentation, has recently received significant attention [16, 21, 23, 27, 36, 42, 49]. This interest is largely attributed to the recent advent of large vision foundation models (*e.g.* SAM [19] or DINOv2 [30]), which are pre-trained with large-scale datasets. By utilizing their well-generalized knowledge, performing segmentation on specific visual concepts, even with fine-tuning a small number of parameters, has become feasible [36, 42, 49]. One straightforward approach for the personalization of semantic segmentation is fine-tuning the text prompts. Despite their outstanding performances, previous studies in the personalization of semantic segmentation fail to incorporate their task with open-vocabulary segmentation. As aforementioned, semantic segmentation models in real-world applications may encounter unseen classes, so personalization of semantic segmentation with novel categories is essential for their applicability. While several methods [2, 7, 10, 36, 49] have shown the capability to segment novel concepts using few-shot samples and masks, they are either not designed for, or limited in, distinguishing personalized concepts among objects within the same class (*e.g.*, identifying ‘my bird’ among ‘multiple birds’). To this end, we propose a simple yet effective plug-in method for personalized OVSS along with our newly established benchmark.

### 3. Method

Before diving into our proposed method, we briefly explain the widely used standard prediction process of OVSS mod-

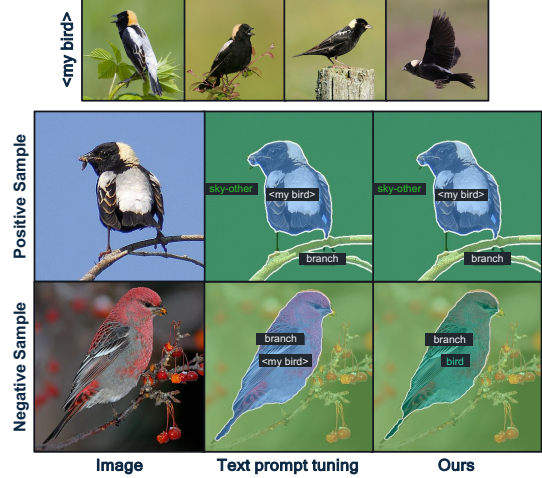


Figure 3. Increased false positives with text prompt tuning. Light-blue pixels and the label ‘<my bird>’ indicate predictions on the personalized concept. The first and second rows show results on positive and negative samples, respectively. In the second row, prompt tuning causes false positives, misclassifying a different bird as ‘<my bird>’ instead of bird. In contrast, our method correctly labels other birds as bird.

els [38, 40]. We have the textual embeddings for open-vocabulary segmentation  $\mathbf{T} \in \mathbb{R}^{V \times D}$ , where  $V$  and  $D$  refer to the open-vocabulary size and the feature dimension, respectively. We also have mask embeddings  $\mathbf{Z} \in \mathbb{R}^{N \times D}$ , where  $N$  indicates the number of masks. We compute the similarity between words and masks  $\mathbf{S} = \mathbf{T} \cdot \mathbf{Z}^T$ , where  $\mathbf{S} \in \mathbb{R}^{V \times N}$ . Then, by using the similarity map  $\mathbf{S}$  and the mask proposals  $\mathbf{M} \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times N}$ , we obtain the final prediction output  $\mathbf{P} = \mathbf{M} \times \mathbf{S}^T$ , where  $\mathbf{P} \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times V}$ . We resize  $\mathbf{P}$  to the original image size for the evaluation.

#### 3.1. Text Prompt Tuning

Fig. 4 describes the overall procedure of our proposed method. Our method first utilizes text prompt tuning to learn the personal visual concept with a few pairs of images and masks. Performing text prompt tuning for fine-tuning the pre-trained vision-language models such as CLIP [32] has been widely used in various multimodal studies [12, 13, 17, 35, 46, 47]. While there are various ways for text prompt tuning, we initialize a learnable embedding vector with textual embedding using the class name of the visual concept we want to personalize. Then, we forward it to the off-the-shelf OVSS model (*e.g.* SAN [40]) along with the original textual embeddings for open-vocabulary segmentation  $\mathbf{T}_{\text{open}} \in \mathbb{R}^{V \times D}$ . In other words, our learnable textual embedding  $\mathbf{T}_{\text{per}} \in \mathbb{R}^{1 \times D}$  is concatenated to  $\mathbf{T}_{\text{open}}$ , leading us to have final textual embeddings  $\mathbf{T} = [\mathbf{T}_{\text{open}}; \mathbf{T}_{\text{per}}]$ . For training the learnable textual embedding, we follow the loss functions used in existing open-vocabulary segmentation models. Specifically, we utilize a dice loss  $\mathcal{L}_{\text{dice}}$ , a binary cross-entropy loss  $\mathcal{L}_{\text{bce}}$  for mask generation, and a cross-entropy loss  $\mathcal{L}_{\text{cls}}$  for mask recognition. The total loss

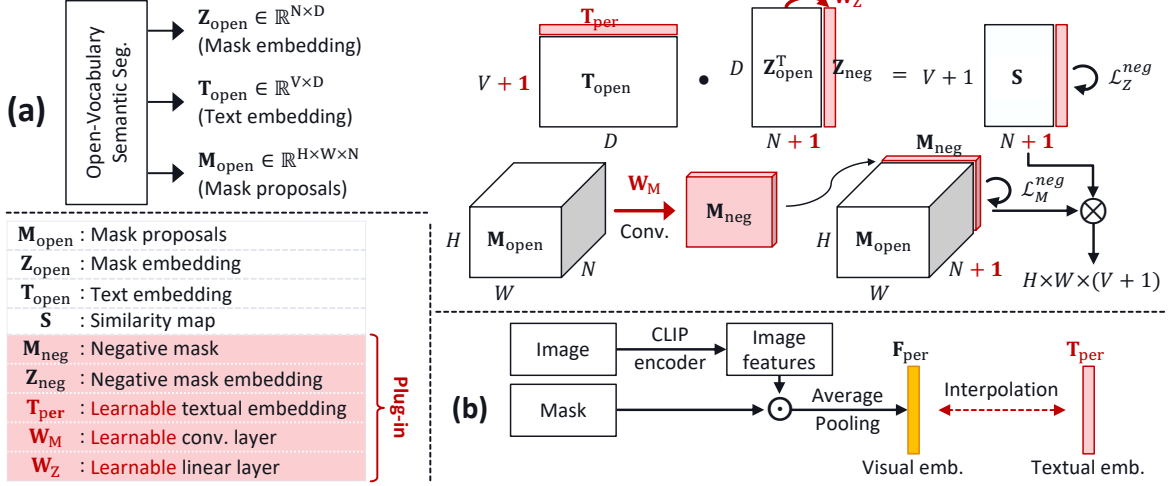


Figure 4. Overview of our proposed method. (a) We additionally use a learnable textual embedding  $T_{per}$  to existing textual embedding  $T_{open}$ , which focuses on learning the personal visual concept. Given existing mask embedding  $Z_{open}$  and mask proposal  $M_{open}$ , we obtain  $Z_{neg}$  and  $M_{neg}$  by forwarding to  $W_Z$  and  $W_M$ , which we add to the original mask embedding and mask proposals, respectively. (b) We obtain a masked visual embedding using an image encoder (e.g. CLIP), and we interpolate it with  $T_{per}$  to enrich the information of personal concept.

is as follows,

$$\mathcal{L}_{seg} = \lambda_1 \mathcal{L}_{dice} + \lambda_2 \mathcal{L}_{bce} + \lambda_3 \mathcal{L}_{cls}, \quad (1)$$

where  $\lambda_1$ ,  $\lambda_2$ , and  $\lambda_3$  are loss weights. We set the loss weights following OVSS models we use [38, 40].

### 3.2. Negative Mask Proposal

While text prompt tuning recognizes the personal visual concept well, we observe that it may lead to increased false positives. Fig. 3 well describes our findings, where we examine whether the visual concept of ‘my bird’ is distinguished among various types of birds via text prompt tuning. When personalizing the OVSS model on the concept of ‘my bird’, text prompt tuning recognizes it well (Fig. 3, first row). However, when an image of a different bird is given as the test sample, the model using text prompt tuning also predicts it as ‘my bird’, indicating that it misrecognizes other similar concepts as the personal visual concept, which leads to increased false positives (Fig. 3, second row). Since one of the goals of this task is to maintain the original OVSS performance, reducing false positives is important. Based on this observation, we propose a negative mask proposal to prevent the model from making false positive predictions.

While not a segmentation work, similar tendency was also observed in recent work, Yo’LLaVA [29], which states that performing personalization with LLaVA [22] using only positive samples leads the model to always answer that the interested subject is in the photo, regardless of the question. While Yo’LLaVA necessitates collecting a non-trivial number of negative samples with the same class to address this issue, our method does not require such a labor-intensive process.

Regarding the mask embeddings, we have the original mask embeddings  $Z_{open} \in \mathbb{R}^{N \times D}$ . Then, we obtain a negative mask embedding  $Z_{neg} \in \mathbb{R}^{N \times D}$  by performing linear combination of  $Z_{open}$  using the learnable weight  $W_Z \in \mathbb{R}^{1 \times N}$ , which is formulated as,

$$Z_{neg} = W_Z Z_{open}. \quad (2)$$

We concatenate  $Z_{neg}$  to  $Z_{open}$  that gives us  $Z = [Z_{open}; Z_{neg}]$ , where  $Z \in \mathbb{R}^{(N+1) \times D}$ .

Then, we obtain the similarity map  $S = T \cdot Z^T$ . Since the goal of  $Z_{neg}$  is to learn words other than the personal visual concept, we supervise it to learn other words equally, which can be formulated using  $S$  as,

$$\mathcal{L}_Z^{neg} = - \sum_{i=1, i \neq k}^V \frac{1}{V-1} \log S[i, j], \quad (3)$$

where  $j$  and  $k$  indicate the index of negative mask embedding and the index of personal concept, respectively.

Given original mask proposals  $M_{open} \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times N}$ , we obtain a mask  $M_{neg} \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times 1}$  that learns visual concepts other than the personal concept, referred to as the negative mask. With a learnable convolutional layer  $W_M$ , we obtain a negative mask by using  $M_{open}$ , which is formulated as,

$$M_{neg} = W_M M_{open}. \quad (4)$$

With the ground truth mask of personal concept  $M_{gt}$ , we use  $1 - M_{gt}$  as the ground truth mask for the supervision of  $M_{neg}$  using the binary cross-entropy loss, which can be formulated as,

$$\mathcal{L}_M^{neg} = -(1 - M_{gt}) \log(M_{neg}) - M_{gt} \log(1 - M_{neg}). \quad (5)$$



For the final mask proposal, we use  $\mathbf{M} = [\mathbf{M}_{\text{open}}; \mathbf{M}_{\text{neg}}] \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times (N+1)}$ .

### 3.3. Injection of Visual Embeddings

Based on the observation that utilizing a single modality (*e.g.* vision or language) for prompt tuning leads to a limited performance gain, recent studies [6, 17, 18] incorporate both representation spaces of vision and language models for prompt tuning. Inspired by such studies, we further improve the performance of personalized OVSS by injecting the visual information into the learnable textual embeddings, as described in Fig. 4 (b). Specifically, we bring a pretrained image encoder  $\mathbf{I}_{\text{enc}}$  (*e.g.* CLIP image encoder) and forward a given image  $\mathbf{X}$ . Then, we extract the feature map  $\mathbf{F} = \mathbf{I}_{\text{enc}}(\mathbf{X})$  and select feature embeddings from the feature map that correspond to the personal concepts using the mask interpolated to match the resolution of  $\mathbf{F}$ , denoted as  $\mathbf{M}'_{\text{gt}}$ , and average them. This masked visual embedding including the personal visual concept,  $\mathbf{F}_{\text{per}}$ , is formulated as,

$$\mathbf{F}_{\text{per}} = \frac{1}{\sum_{j=1}^{HW} \mathbb{1}(\mathbf{M}'_{\text{gt}} = 1)} \sum_{i=1}^{HW} \mathbf{F} \odot \mathbf{M}'_{\text{gt}}, \quad (6)$$

where  $\odot$  indicates the element-wise multiplication. When multiple images are given, we average  $\mathbf{F}_{\text{per}}$  of multiple images. Finally, during text prompt tuning, we replace  $\mathbf{T}_{\text{per}}$  with  $\mathbf{T}_{\text{per}}^{\text{vis}} = \alpha \cdot \mathbf{F}_{\text{per}} + (1 - \alpha) \cdot \mathbf{T}_{\text{per}}$ , which is the interpolation between the masked visual embedding  $\mathbf{F}_{\text{per}}$  and the personalized textual embedding  $\mathbf{T}_{\text{per}}$ . To this end, our total loss function is formulated as,

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{seg}} + \lambda_Z^{\text{neg}} \mathcal{L}_Z^{\text{neg}} + \lambda_M^{\text{neg}} \mathcal{L}_M^{\text{neg}}. \quad (7)$$

## 4. Experiments

### 4.1. Experimental Settings

**Datasets.** We conduct experiments using three datasets: 1) FSS-1000 [21], 2) CUB-200 [34], and 3) ADE-20K-847 [44]. Specifically, we bring pretrained OVSS models and perform the personalization of a certain visual concept on these three datasets. The models are pretrained with COCO Stuff [4], so we assume that classes included in the three datasets used for evaluation are regarded as open-vocabulary words.

We evaluate the models depending on the number of images  $K$  used for personalization, which we set to  $K = 1, 3, 5$ . For the test sets, we include equal number of images including the personal concept and those without the personal concept. We intentionally include the images without the personal concept since models that are fine-tuned to recognize personal concepts overconfidently predict other concepts as personal concepts as well. Preventing such behavior is one of the main goals of our work, so we include

images without personal concepts in our test sets. We select 30 classes for FSS-1000 and ADE-20K-847, while using all 200 classes for CUB-200. Since the selected classes and evaluation settings are different from the original datasets, we refer to the processed versions of FSS-1000, CUB-200, and ADE-20K as  $\text{FSS}^{\text{per}}$ ,  $\text{CUB}^{\text{per}}$ , and  $\text{ADE}^{\text{per}}$ , respectively. Further details are included in the supplementary.

**Evaluation Metrics.** We mainly utilize two different evaluation metrics using intersection over union (IoU):  $\text{IoU}^{\text{per}}$  and  $\text{mIoU}$ .  $\text{IoU}^{\text{per}}$  evaluates how precisely a given segmentation model predicts the personal concept. Specifically, the regions of the personal concept are considered as the positive labels, while the regions of other classes are all negative labels. For  $\text{mIoU}$ , in addition to  $\text{IoU}^{\text{per}}$ , we include IoU of other classes and average the values of IoU. Since ADE-20K-847 originally includes the ground truth labels for open-vocabulary classes, we use those labels for evaluating  $\text{mIoU}$  in  $\text{ADE}^{\text{per}}$ . However, FSS-1000 and CUB-200 do not include the ground truth labels for open-vocabulary classes, so we assume the predictions of pretrained open-vocabulary segmentation models as the ground truth labels for calculating  $\text{mIoU}$  for  $\text{FSS}^{\text{per}}$  and  $\text{CUB}^{\text{per}}$ . The main goal of this work is to improve  $\text{IoU}^{\text{per}}$  while maintaining a reasonable level of performance on  $\text{mIoU}$ .

**Implementation Details.** For the pretrained open-vocabulary segmentation models, we utilize two representative open-vocabulary segmentation models, Side Adapter Network (SAN) [40] and Open-vocabulary Diffusion-based panoptic SEgmentation (ODISE) [38]. For SAN and ODISE, we use ViT-B/16 CLIP model and ViT-L/14 CLIP model, respectively. Since ODISE utilizes the diffusion model, we use the Stable Diffusion v1.3 model [33] for it. We train  $\mathbf{T}_{\text{per}}$ ,  $\mathbf{W}_M$ , and  $\mathbf{W}_Z$  for 200 iterations. We initialize the learnable textual embedding with the textual embedding of the label of the personal concept using the text description ‘a photo of’ before the label. Further details on our experimental settings are included in the supplementary.

### 4.2. Quantitative Results

For quantitative evaluation, we compare open-vocabulary segmentation models without personalization and those with our method. Specifically, the models without personalization employ only text prompts for recognizing personal visual concept, such as ‘black footed albatross’. Table 1 compares  $\text{IoU}^{\text{per}}$  and  $\text{mIoU}$  in personalized OVSS. Across three datasets, our method improves  $\text{IoU}^{\text{per}}$  when applied to both SAN and ODISE, demonstrating its scalability. For example, by average, we improve  $\text{IoU}^{\text{per}}$  of SAN by 12.48, 8.65, 15.79 for  $\text{FSS}^{\text{per}}$ ,  $\text{CUB}^{\text{per}}$ , and  $\text{ADE}^{\text{per}}$ , respectively. While our method requires a few pairs of images and masks for personal concept, we still improve  $\text{IoU}^{\text{per}}$  even with one sample, verifying the practicality in real-world applications.

Along with improving  $\text{IoU}^{\text{per}}$ , maintaining the origi-

Dataset	Method	IoU <sup>per</sup>				mIoU			
		$K = 1$	$K = 3$	$K = 5$	Avg.	$K = 1$	$K = 3$	$K = 5$	Avg.
FSS <sup>per</sup>	ODISE [38]	10.69	10.69	10.69	10.69	<b>23.86</b>	<b>23.86</b>	<b>23.86</b>	<b>23.86</b>
	+ Ours	<b>30.97</b>	<b>33.11</b>	<b>34.05</b>	<b>32.71</b>	21.83	22.68	22.94	22.48
	SAN [40]	41.08	41.08	41.08	41.08	55.68	55.68	55.68	55.68
	+ Ours	<b>49.80</b>	<b>54.09</b>	<b>56.80</b>	<b>53.56</b>	<b>56.40</b>	<b>56.73</b>	<b>55.85</b>	<b>56.32</b>
CUB <sup>per</sup>	ODISE [38]	0.02	0.02	0.02	0.02	<b>47.48</b>	<b>47.48</b>	<b>47.48</b>	<b>47.48</b>
	+ Ours	<b>5.39</b>	<b>5.90</b>	<b>5.99</b>	<b>5.76</b>	45.16	44.94	44.88	44.99
	SAN [40]	68.25	68.25	68.25	68.25	77.32	77.32	77.32	77.32
	+ Ours	<b>76.70</b>	<b>77.21</b>	<b>76.80</b>	<b>76.90</b>	<b>77.36</b>	<b>77.85</b>	<b>78.29</b>	<b>77.83</b>
ADE <sup>per</sup>	ODISE [38]	1.24	1.24	1.24	1.24	<b>12.22</b>	<b>12.22</b>	<b>12.22</b>	<b>12.22</b>
	+ Ours	<b>9.45</b>	<b>10.85</b>	<b>13.43</b>	<b>11.24</b>	12.21	12.18	12.18	12.19
	SAN [40]	6.88	6.88	6.88	6.88	17.20	17.20	<b>17.20</b>	17.20
	+ Ours	<b>17.08</b>	<b>24.80</b>	<b>26.15</b>	<b>22.67</b>	<b>17.26</b>	<b>17.21</b>	17.19	<b>17.22</b>

Table 1. Comparisons on IoU<sup>per</sup> and mIoU. We apply our method on both SAN [40] and ODISE [38] on FSS<sup>per</sup>, CUB<sup>per</sup>, and ADE<sup>per</sup>. We vary  $K$ , the number of images and masks, to 1, 3, and 5.

Text Prompt	Neg. Mask	Visual Inject	mIoU	IoU <sup>per</sup>	IoU <sup>per</sup> <sub>precision</sub>	IoU <sup>per</sup> <sub>recall</sub>
-	-	-	77.32	68.25	92.25	72.95
✓	-	-	77.89	69.70	74.75	91.04
✓	✓	-	77.89	73.71	80.07	90.17
✓	-	✓	77.65	65.94	70.06	91.58
✓	✓	✓	<b>78.29</b>	<b>76.80</b>	84.51	89.07

Table 2. Ablation study on our proposed method. We show how each module contributes to the performance gain by reporting mIoU, IoU<sup>per</sup>, IoU<sup>per</sup><sub>precision</sub> and IoU<sup>per</sup><sub>recall</sub>.

nal OVSS performance is also important for personalized OVSS. Table 1 also shows that our method maintains mIoU for both SAN [40] and ODISE [38] across three datasets. This indicates that our method does not compromise the OVSS performance for recognizing personal visual concepts. Maintaining the original OVSS performance is crucial for deploying such models in real-world applications since users do not want their models to misrecognize objects that the models used to recognize well. To the best of our knowledge, this is the first work to establish the experimental setting of personalized OVSS task. We believe that establishing such an experimental setting itself serves as a large contribution to this field. Additionally, since this is a newly proposed task, we lack baseline methods beyond the open-vocabulary segmentation models without personalization. For further demonstrating the superiority of our proposed method, we compare our method with an approach that replaces the learnable text prompts with visual embeddings, which we report in our supplementary.

## 5. Further Analysis

**Ablation Study** For an in-depth analysis of our proposed method, we conduct experiments to investigate how each component of our proposed method contributes to the performance gain. Specifically, we show how applying text

prompt tuning, utilizing negative mask proposal, and injecting visual embedding improves the performance. For ablation study, we use CUB<sup>per</sup> with  $K = 5$  and SAN for the dataset and the pretrained model, respectively. In addition to mIoU and IoU<sup>per</sup>, we report IoU<sup>per</sup><sub>precision</sub> and IoU<sup>per</sup><sub>recall</sub> to analyze how IoU<sup>per</sup> improves regarding true and false positive predictions.

As shown in Table 2, we observe that applying text prompt tuning improves recognizing the personal visual concept (IoU<sup>per</sup> of 69.70) compared to the one without personalization (IoU<sup>per</sup> of 68.25). While text prompt tuning improves IoU<sup>per</sup> by increasing IoU<sup>per</sup><sub>recall</sub>, we observe a large performance degradation on IoU<sup>per</sup><sub>precision</sub>, indicating significantly increased false positives since the model overconfidently predicts other objects as the personal visual concept. To address this issue, we utilize the negative mask proposal that improves IoU<sup>per</sup><sub>precision</sub> from 74.75 to 80.07, improving IoU<sup>per</sup> to 73.71. On the other hand, while combining text prompt tuning with visual injection improves IoU<sup>per</sup><sub>recall</sub>, it rather degrades IoU<sup>per</sup><sub>precision</sub>. Our final method achieves the best performance (IoU<sup>per</sup> of 76.80) with reasonable level of both IoU<sup>per</sup><sub>precision</sub> and IoU<sup>per</sup><sub>recall</sub>.

**Segmentation Results** Fig. 5 and Fig. 6 shows the segmentation results of our method on several samples from COCO dataset [4] and CUB<sup>per</sup>, respectively. Fig. 5 shows that personalized OVSS enables to identify my dog among multiple dogs while OVSS without personalization fails to differentiate my dog with other dogs. Fig. 6 visualizes the effect of each module in our method. As shown in the first row of Fig. 6, when performing personalization with only visual embedding injected, the model also captures the other bird as my bird, showing increased false positives. Additionally, when only the negative mask proposal is applied without the visual embedding injected, it fails to capture my bird. Therefore, we need both modules to capture my bird while




Figure 5. Qualitative results of personalized OVSS on test samples from COCO dataset. While existing OVSS models simply recognize different dogs as a dog, personalized OVSS identifies ‘my dog’ among two dogs in the same image.

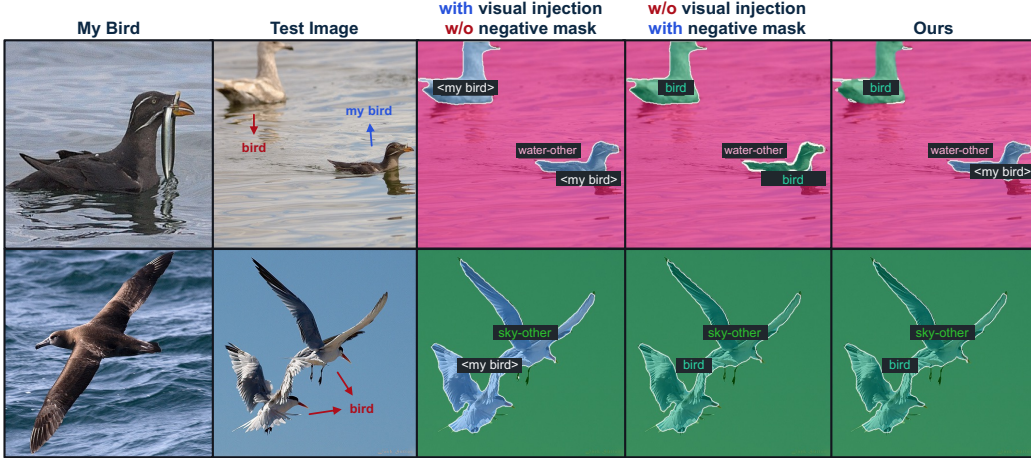


Figure 6. Segmentation results on CUB-200. We show the effect of each module in our method. Visual embedding injection alone identifies my bird well but increases false positives. Negative mask alone reduces false positives but may miss my bird. Therefore, combining both is essential for accurate segmentation.

not recognizing other birds as my bird. The second row also shows similar tendency when there is a photo without my bird. Such a qualitative analysis clearly demonstrates that utilizing negative mask proposal prevents the model from making false positives and visual embedding injection enriches the representation of textual embedding towards the personal concept.

**Comparison with Related Approaches** To further demonstrate that our proposed personalized OVSS is a challenging task, we compare our method with recent 1) open-vocabulary method [10] and 2) few-shot segmentation methods [36, 49]. Tab. 3 demonstrates that kNN-CLIP [10] using CLIP fails to recognize the personal concept. Also, Tab 4 compares our method with SEEM [49] and SegGPT [36], which are recent state-of-the-art few-shot segmentation methods. As shown, while SEEM and SegGPT capture birds well (improved  $\text{IoU}_{\text{recall}}^{\text{per}}$ ), they recognize other birds also as the bird we want to personalize (large degra-

dation in  $\text{IoU}_{\text{precision}}^{\text{per}}$ ). Such results show that even recent models addressing similar tasks struggle to capture my personal concept among objects of same class.

Dataset	Metric	$\text{IoU}^{\text{per}}$	mIoU
CUB <sup>per</sup> ( $K = 5$ )	kNN-CLIP	1.27	56.34
	Ours	<b>76.80</b>	<b>78.29</b>

Table 3. Comparisons with kNN-CLIP on CUB<sup>per</sup> with  $K = 5$ .

Dataset	Methods	$\text{IoU}^{\text{per}}$	$\text{IoU}_{\text{precision}}^{\text{per}}$	$\text{IoU}_{\text{recall}}^{\text{per}}$
CUB <sup>per</sup> ( $K = 1$ )	SEEM	45.54	47.64	90.82
	SegGPT	47.00	48.56	<b>93.88</b>
	Ours	<b>76.80</b>	<b>84.51</b>	89.07

Table 4. Comparisons with recent few-shot semantic segmentation methods on CUB<sup>per</sup> with  $K = 1$ .






Figure 7. Visualization of negative mask proposal. The white and black pixels in the ground truth masks indicate the personal visual concept and other regions, respectively. Our negative mask proposal predicts the opposite, indicating that it learned the concepts other than the personal concept.

### 5.1. Visualization of Negative Masks

We also visualize the negative mask proposals, which we propose in order to prevent the model from making overconfident predictions on other classes when learning the personal visual concept. As aforementioned, newly attached negative mask focuses on learning all concepts other than the personal visual concept. In other words, it outputs low prediction values for regions associated with the personal visual concept and high values elsewhere. Fig. 7 shows that the negative mask well recognizes regions other than the personal visual concept. Specifically, white and black pixels in ground truth masks indicate the personal visual concept and other concepts, respectively. The negative mask proposal outputs the opposite values compared to the ground truth mask, indicating that it outputs low prediction values for the personal visual concept and high values for other concepts. Adding such a mask focusing on other concepts prevents the model from making false positives, as demonstrated with improved  $\text{IoU}_{\text{precision}}^{\text{per}}$  in Table 2.

### 5.2. Discussion for Personalized OVSS

Personalized open-vocabulary semantic segmentation differs from traditional few-shot segmentation [11, 15, 24, 36, 41] or class-incremental learning [26, 28, 45, 48]. Unlike few-shot segmentation that focuses only on recognizing a specific given concept, personalized open-vocabulary semantic segmentation aims to maintain the capability of open-vocabulary segmentation to segment objects using arbitrary text descriptions while recognizing new personal visual concepts. To enable various open-vocabulary segmentation models to understand personalized concepts, we design a simple and effective plug-in method based on text prompt tuning. Furthermore, different from class-incremental learning, which learns new classes sequentially while retaining previous ones without open-vocabulary capability, our method leverages text prompt tuning to capture personal concept based on open-vocabulary capability.

This approach enables segmentation of specific objects

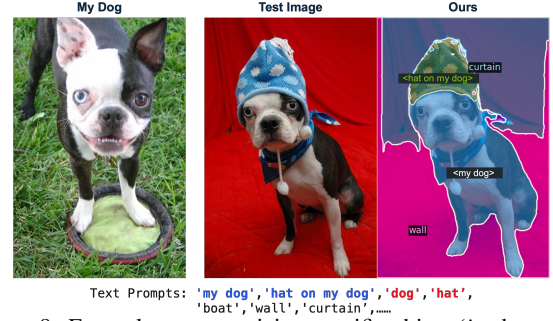


Figure 8. Example on recognizing specific object (*i.e.* hat on my dog) by combining the learned personalized prompt (*i.e.* my dog) with the text prompt (*i.e.* hat on). Despite the presence of confusing text prompts such as ‘dog’ and ‘hat’, it can be observed that the model recognizes the text prompts ‘my dog’ and ‘hat on my dog’ when including the learned personalized prompts.

by combining text descriptions with personalized prompts. For example, our method can recognize both ‘my dog’ and ‘hat on my dog’ instead of ‘dog’ and ‘hat’, as shown in Fig. 8. We believe that the use cases shown in Fig. 8 clearly demonstrates the necessity of our newly proposed task since the current studies in OVSS, few-shot learning, and class-incremental learning fail to recognize ‘hat on my dog’. Although we did not include such experiments in our quantitative evaluation due to the absence of the evaluation dataset, we believe that extending our work in such a direction would be an interesting future work.

## 6. Conclusions

In conclusion, this study proposes a novel task termed *personalized open-vocabulary semantic segmentation*, which captures the personal visual concept provided from few images and masks utilizing a pretrained OVSS model. Unlike existing few-shot segmentation or OVSS, which are limited to segmenting only given visual concepts or fail to understand personalized vocabulary, we extend OVSS models to also comprehend personalized vocabulary, enabling them to identify the personal concept among objects within the same category. As aforementioned, such a task is important when asking my robot assistant to find my favorite mug cup among multiple mug cups in a cabinet. Observing that existing studies do not have an established benchmark for this, we conduct experiments on our newly proposed experimental setting with processed datasets  $\text{FSS}^{\text{per}}$ ,  $\text{CUB}^{\text{per}}$ , and  $\text{ADE}^{\text{per}}$ . This paper also proposes a method including text prompt tuning, negative mask proposal, and injection of visual embedding, which could be applied to existing OVSS methods. Considering the importance of this task, we hope that our work will inspire future researchers to delve into personalized AI systems from recognizing only general words to also understanding personal words, thus moving towards more personalized AI that adapts to individual contexts and preferences.



**Acknowledgement** We would like to thank the Qualcomm AI Research team for their valuable discussions.

## References

- [1] Donghyeon Baek, Youngmin Oh, and Bumsub Ham. Exploiting a joint embedding space for generalized zero-shot semantic segmentation. In *ICCV*, 2021. 3
- [2] Ivana Balazevic, David Steiner, Nikhil Parthasarathy, Relja Arandjelovic, and Olivier J Henaff. Towards in-context scene understanding. In *NeurIPS*, 2023. 3
- [3] Maxime Bucher, Tuan-Hung VU, Matthieu Cord, and Patrick Pérez. Zero-shot semantic segmentation. In *NeurIPS*, 2019. 3
- [4] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. Coco-stuff: Thing and stuff classes in context. In *CVPR*, 2018. 5, 6
- [5] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018. 1
- [6] Eulrang Cho, Jooyeon Kim, and Hyunwoo J Kim. Distribution-aware prompt tuning for vision-language models. In *ICCV*, 2023. 5
- [7] Niv Cohen, Rinon Gal, Eli A Meirom, Gal Chechik, and Yuval Atzmon. “this is my unicorn, fluffy”: Personalizing frozen vision-language representations. In *European conference on computer vision*, pages 558–577. Springer, 2022. 3
- [8] Jian Ding, Nan Xue, Gui-Song Xia, and Dengxin Dai. Decoupling zero-shot semantic segmentation. In *CVPR*, 2022. 2, 3
- [9] Zhangxuan Gu, Siyuan Zhou, Li Niu, Zihan Zhao, and Liqing Zhang. Context-aware feature generation for zero-shot semantic segmentation. In *ACM International Conference on Multimedia*, 2020. 3
- [10] Zhongrui Gui, Shuyang Sun, Runjia Li, Jianhao Yuan, Zhaochong An, Karsten Roth, Ameya Prabhu, and Philip Torr. knn-clip: Retrieval enables training-free segmentation on continually expanding large vocabularies. *CoRR*, 2024. 3, 7
- [11] Sina Hajimiri, Malik Boudiaf, Ismail Ben Ayed, and Jose Dolz. A strong baseline for generalized few-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11269–11278, 2023. 2, 8
- [12] Changsheng Xu Hantao Yao, Rui Zhang. Tcp: Textual-based class-aware prompt tuning for visual-language model. In *CVPR*, 2024. 3
- [13] Siteng Huang, Biao Gong, Yulin Pan, Jianwen Jiang, Yiliang Lv, Yuyuan Li, and Donglin Wang. VoP: Text-video co-operative prompt tuning for cross-modal retrieval. In *CVPR*, 2023. 3
- [14] D. Huynh, J. Kuen, Z. Lin, J. Gu, and E. Elhamifar. Open-vocabulary instance segmentation via robust cross-modal pseudo-labeling. In *CVPR*, 2022. 3
- [15] Lai Jinxiang, Yang Siqian, Liu Wenlong, Zeng Yi, Huang Zhongyi, Wu Wenlong, Liu Jun, Gao Bin-Bin, and Wang Chengjie. tsf: Transformer-based semantic filter for few-shot learning. In *ECCV*, 2022. 2, 8
- [16] Aliasghar Khani, Saeid Asgari, Aditya Sanghi, Ali Mahdavi Amiri, and Ghassan Hamarneh. SLime: Segment like me. In *ICLR*, 2024. 3
- [17] Muhammad Uzair khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Maple: Multi-modal prompt learning. In *CVPR*, 2023. 3, 5
- [18] Muhammad Uzair Khattak, Syed Talal Wasim, Muzammal Naseer, Salman Khan, Ming-Hsuan Yang, and Fahad Shahbaz Khan. Self-regulating prompts: Foundational model adaptation without forgetting. In *ICCV*, 2023. 5
- [19] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollar, and Ross Girshick. Segment anything. In *ICCV*, 2023. 3
- [20] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and Rene Ranftl. Language-driven semantic segmentation. In *ICLR*, 2022. 3
- [21] Xiang Li, Tianhan Wei, Yau Pun Chen, Yu-Wing Tai, and Chi-Keung Tang. Fss-1000: A 1000-class dataset for few-shot segmentation. In *CVPR*, 2020. 3, 5
- [22] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning, 2023. 4
- [23] Yang Liu, Muzhi Zhu, Hengtao Li, Hao Chen, Xinlong Wang, and Chunhua Shen. Matcher: Segment anything with one shot using all-purpose feature matching. In *ICLR*, 2024. 3
- [24] Zhihe Lu, Sen He, Xiatian Zhu, Li Zhang, Yi-Zhe Song, and Tao Xiang. Simpler is better: Few-shot semantic segmentation with classifier weight transformer. In *ICCV*, 2021. 2, 8
- [25] Huaishao Luo, Junwei Bao, Youzheng Wu, Xiaodong He, and Tianrui Li. SegCLIP: Patch aggregation with learnable centers for open-vocabulary semantic segmentation. In *ICML*, 2023. 2
- [26] Marc Masana, Xialei Liu, Bartłomiej Twardowski, Mikel Menta, Andrew D. Bagdanov, and Joost van de Weijer. Class-incremental learning: Survey and performance evaluation on image classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. 8
- [27] Lingchen Meng, Shiyi Lan, Hengduo Li, Jose M. Alvarez, Zuxuan Wu, and Yu-Gang Jiang. Segic: Unleashing the emergent correspondence for in-context segmentation, 2024. 3
- [28] Sudhanshu Mittal, Silvio Galesso, and Thomas Brox. Essentials for class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2021. 8
- [29] Thao Nguyen, Haotian Liu, Yuheng Li, Mu Cai, Utkarsh Ojha, and Yong Jae Lee. Yo’llava: Your personalized language and vision assistant, 2024. 2, 4
- [30] Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran, Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Herve Jegou,

- Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision, 2023. 3
- [31] Jie Qin, Jie Wu, Pengxiang Yan, Ming Li, Ren Yuxi, Xuefeng Xiao, Yitong Wang, Rui Wang, Shilei Wen, Xin Pan, et al. Freeseg: Unified, universal and open-vocabulary image segmentation. In *CVPR*, 2023. 3
- [32] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 1, 2, 3
- [33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 3, 5
- [34] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. Caltech-ucsd birds 200. Technical report, California Institute of Technology, 2011. 5
- [35] Fengxiang Wang, Wanrong Huang, Shaowu Yang, Qi Fan, and Long Lan. Learning to learn better visual prompts. 2024. 3
- [36] Xinlong Wang, Xiaosong Zhang, Yue Cao, Wen Wang, Chunhua Shen, and Tiejun Huang. Seggpt: Segmenting everything in context. In *ICCV*, 2023. 3, 7, 8
- [37] Yongqin Xian, Subhabrata Choudhury, Yang He, Bernt Schiele, and Zeynep Akata. Semantic projection network for zero- and few-label semantic segmentation. In *CVPR*, 2019. 1, 2, 3
- [38] Jiarui Xu, Sifei Liu, Arash Vahdat, Wonmin Byeon, Xiaolong Wang, and Shalini De Mello. Open-vocabulary panoptic segmentation with text-to-image diffusion models. In *CVPR*, 2023. 1, 3, 4, 5, 6
- [39] Mengde Xu, Zheng Zhang, Fangyun Wei, Yutong Lin, Yue Cao, Han Hu, and Xiang Bai. A simple baseline for open vocabulary semantic segmentation with pre-trained vision-language model. In *ECCV*, 2022.
- [40] Mengde Xu, Zheng Zhang, Fangyun Wei, Han Hu, and Xiang Bai. Side adapter network for open-vocabulary semantic segmentation. In *CVPR*, 2023. 2, 3, 4, 5, 6
- [41] Gengwei Zhang, Guoliang Kang, Yi Yang, and Yunchao Wei. Few-shot segmentation via cycle-consistent transformer. In *NeurIPS*, 2021. 8
- [42] Renrui Zhang, Zhengkai Jiang, Ziyu Guo, Shilin Yan, Junting Pan, Hao Dong, Yu Qiao, Peng Gao, and Hongsheng Li. Personalize segment anything model with one shot. In *ICLR*, 2024. 3
- [43] Zhuowen Tu Zheng Ding, Jieke Wang. Open-vocabulary universal image segmentation with maskclip. In *ICML*, 2023. 2, 3
- [44] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *CVPR*, 2017. 5
- [45] Da-Wei Zhou, Qi-Wei Wang, Zhi-Hong Qi, Han-Jia Ye, De-Chuan Zhan, and Ziwei Liu. Class-incremental learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 8
- [46] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *CVPR*, 2022. 3
- [47] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *IJCV*, 2022. 3
- [48] Fei Zhu, Zhen Cheng, Xu-yao Zhang, and Cheng-lin Liu. Class-incremental learning via dual augmentation. In *NeurIPS*, 2021. 8
- [49] Xueyan Zou, Jianwei Yang, Hao Zhang, Feng Li, Linjie Li, Jianfeng Wang, Lijuan Wang, Jianfeng Gao, and Yong Jae Lee. Segment everything everywhere all at once. In *NeurIPS*, 2023. 3, 7